

A Comparison of Techniques for Scheduling Fleets of Earth-Observing Satellites

Al Globus
CSC
NASA Ames

James Crawford
RIACS
NASA Ames

Jeremy Frank
NASA Ames

Jason Lohn
NASA Ames

Anna Pryor
NASA Ames

David Smith
NASA Ames

Abstract

We have run realistically-sized experiments comparing a number of techniques for the Earth-Observing Satellite (EOS) scheduling problem. This problem requires the scheduler to satisfy complex constraints while attempting to include as many high-priority observations as possible. We compared the genetic algorithm, simulated annealing, squeaky wheel optimization, stochastic hill climbing, iterated sampling, and heuristic-based stochastic search with contention heuristics under a number of conditions on a three satellite, one week, 6,000+ observation problem. Simulated annealing with simple mutation operators was the clear winner. Heuristic-based stochastic search with contention heuristics was hundreds of times slower than the simpler techniques and performed very poorly even when far more CPU was used. For this problem, at least among the techniques tested, simple fast and stupid with a little learning significantly out-performs smart without learning.

Introduction

A growing fleet of scientific, military, and commercial Earth observing satellites (EOS) circles the globe. Most of these satellites are within about 700 km of the surface. Observations of any particular location can only be made only when the satellite is overhead. A single orbit takes approximately 100 minutes and, since the Earth turns underneath the satellite, any particular point on the surface is only occasionally, if predictably, visible. Although there are approximately 60 EOS satellites in orbit today, image collection is nearly always scheduled separately for each satellite with manual coordination, if any. Various studies (Globus *et al.* 2002) (Rao, Soma, & Padmashree 1998) have suggested that automatic coordination of multiple satellites can be beneficial, but the best techniques to use is not clear.

This study compares 14 EOS scheduling techniques on a realistic-sized model problem. In particular, we compare simulated annealing, hill climbing, the genetic algorithm, squeaky wheel optimization, iterated sampling (ISAMP) and heuristic based stochastic search (HBSS) with contention heuristics (Frank *et al.* 2002). In the next section we describe the scheduling problem and our model. A description of the scheduling

techniques follows. The nature and results of our computational experiments are then presented along with analysis. Space limitations preclude a review of previous work, but (Globus *et al.* 2002) and (Globus *et al.* 2003) provide such a discussion.

EOS Scheduling Problem

In this section we first describe the EOS scheduling problem as perceived by satellite operators and developers. Then we describe our general model of the problem, with many details removed, and the particular model used to generate our results.

EOS scheduling attempts to schedule as many highest-priority observations as possible within a fixed period of time on a fixed set of satellite born sensors. For example, the Landsat 7 satellite scheduler is considered to have done a good job if 250 observations are made each day, out of perhaps 300 or more requests. EOS is complicated by a number of important constraints. Potin (Potin 1998) lists some of these constraints as:

1. Revisit limitations. A target must be within sight of the satellite; and EOS satellites travel in fixed orbits. These orbits pass over any particular place on Earth at limited times so there are only a few observation windows (and sometimes none) for a given target.
2. Time required to take each image. Most Earth observing satellites take a one dimensional image and use the spacecraft's orbital motion to sweep out the area to be imaged.
3. Limited on-board data storage. Images are typically stored on a solid state recorder (SSR) until they can be sent to the ground.
4. Ground station availability (SSR dumps). The data in the SSR is sent to the ground when the satellite passes over a ground station. Ground station windows are as limited as with any other target.
5. Transition time between look angles (slewing). Some instruments are mounted on motors that can point side-to-side (cross-track).
6. Power and thermal control.

7. Coordination of multiple satellites.
8. Cloud cover. Some sensors cannot see through clouds. Not only do clouds cover much of the Earth at any given time, but some locations are nearly always cloudy.
9. Stereo pair acquisition or multiple observations of the same target by different sensors or the same sensor at different times.

For further details of the EOS scheduling problem see (Frank *et al.* 2002) and (Sherwood *et al.* 1998).

Our model problem implements all these constraints except the last two. The model problem consists of three satellites in Sun-synchronous orbits (orbits in which the equator is crossed at the same local time each orbit) for one week. The satellites are spaced ten minutes apart. Each satellite carries one sensor mounted on a cross-track slewable motor that can point up to 24 degrees side to side and turns one degree in two seconds. Each satellite has an SSR capable of storing 50 arbitrary units.

We model our power and thermal constraints using so called duty cycle constraints, the approach taken by NASA's Landsat 7 satellite. A duty cycle constraint requires that the sensor not be turned on for longer than a maximum time within any interval of a certain length. This insures conformance with power, thermal, and other physical constraints on the space craft. Our model problem uses the Landsat 7 duty cycles. Specifically, an sensor may not be used for more than:

1. 34 minutes in any 100 minute period.
2. 52 minutes in any 200 minute period.
3. 131 minutes in any 600 minute period.

There is one ground station in Alaska. Whenever a satellite comes within sight of the ground station it is assumed to completely empty its SSR, which is then available for additional observation storage. There are approximately 75 SSR dumps per spacecraft during the week. Since some orbits are over oceans and all targets are on land, some SSR dump opportunities are wasted on an empty SSR.

6300 observation targets were randomly generated on land. Of these, 6114 were observable by at least one satellite. The targets were assumed to be in the center of a rectangle that requires 24 seconds (taken from Landsat) of satellite motion to image. Observations require one, three, or five arbitrary storage units (one third each) on the SSR. Each observation is assigned a priority from one to six evenly spaced in 0.1 increments. Each observation has 2-24 windows, times when at least one satellite is within view of the observation's target.

The fitness (quality) of each schedule is determined by a weighted sum (smaller numbers indicate better fitness):

$$F = w_p \sum_{O_u} P_o + w_s M_s + w_{point} M_p \quad (1)$$

where F is the fitness, O_u is the set of unscheduled observation, P_o is the priority of an observation, M_s is the mean time spent slewing per scheduled observation, M_p is the mean off-nadir (nadir is straight down) pointing angle for scheduled observations, w stands for weight, $w_p = 1$, $w_s = 0.01$, and $w_{point} = 0.00137$. Note that the weights heavily favor the priority of unscheduled observations over pointing and slewing time objectives.

Scheduling Techniques

This study compares 14 search techniques applied to the EOS scheduling problem. The simplest techniques tested here are simulated annealing, hill climbing, two variants of the genetic algorithm, and ISAMP (essentially random scheduling) taking random steps. By using a more intelligent mutation operator, these algorithms become variants of squeaky wheel optimization (Joslin & Clements 1999). Finally, we examine HBSS (Bresina 1996) with contention heuristics (Frank *et al.* 2002) where a great deal of processing is devoted to determining the order in which observations are placed in schedule timelines.

We represent a schedule as a permutation (the genotype) of the observations. A simple, deterministic greedy scheduler assigns resources to the observations in the order indicated by the permutation (except for HBSS). This produces a timeline (the phenotype) with all of the scheduled observations, the time they are executed, and the resources (SSR, sensor, pointing angle) used. The greedy scheduler assigns times and resources to observations using earliest-first scheduling heuristics while maintaining consistency with sensor availability, onboard memory (SSR) and slewing constraints. If an observation cannot be scheduled without violating the current constraints (those created by scheduling observations from earlier in the permutation), the observation is left unscheduled.

Simple earliest-first scheduling starting at epoch (time = 0) had some problems. We discovered that the algorithm works better if 'earliest-first' starts at some random. If the observation cannot be scheduled before the end of time, the algorithm starts at epoch and continues until the observation is scheduled or the initial time is reached. The time each observation is scheduled is stored along with the permutation and is preserved by mutation and crossover and is used as the starting point for the earliest-first scheduler operating on modified permutation. The extra scheduling flexibility may explain why this approach works better than earliest-first starting at epoch.

Constraints are enforced by representing sensors and satellite-SSR as timelines. Scheduling an observation causes timelines to take on appropriate values (i.e., in use for a sensor, slew motor setting, amount of SSR memory available) at different times. An observation is inserted at the first time available consistent with the current constraints.

The simplest algorithm tested was ISAMP, which is essentially a random search. With ISAMP, each sched-

ule is generated from a new random permutation with random start times for the simple, deterministic, greedy scheduler.

The next class of algorithms tested were the 'evolutionary' search techniques, which we define here as those that start with random permutations and generate new permutations with mutation and/or crossover. These search techniques are slightly more intelligent than ISAMP in that they remember at least one previously generated schedule and thereby learn. The algorithms tested were:

1. Stochastic hill climbing.
2. Simulated annealing. The temperature starts at 100 (arbitrary units) and is multiplied by 0.92 every 1000 children (100,000 children are generated per run).
3. A steady-state tournament selection genetic algorithm with population size 100. The individual to replace is chosen by a tournament from the whole population where the least fit is replaced. Tournament size is always two.
4. A generational elitist genetic algorithm. The population size is 110 and the 10 best individuals are copied into the next generation.

Each search technique was tested with three mutation operators:

1. Random swap. Two permutation locations are chosen at random and the observations are swapped, with 1-15 swaps (chosen at random) per mutation. Earlier experiments determined that allowing more than one swap improved scheduling. A single random swap is called order-based mutation (Syswerda & Palmucci 1991).
2. Temperature-dependent swap. Here the number of swaps (1-15) is still chosen at random but with a bias. Early in evolution a larger number of swaps tend to be used, and later in evolution fewer swaps are performed. This is analogous to the 'temperature' dependent behavior of simulated annealing. The choice of the number of swaps is determined by a weighted roulette wheel where the weights vary linearly as evolution proceeds starting at n and ending at $15 - n$ where n is the number of swaps. Earlier experiments tried fewer swaps early in evolution and more swaps later. This didn't work as well.
3. Squeaky shift. This mutator implements squeaky wheel optimization. The mutator shifts 1-15 (chosen randomly) 'deserving' observations earlier in the permutation. Early in the permutation an observation is more likely to be scheduled since fewer other observations will have been scheduled creating additional constraints. Each observation to shift forward is chosen by a tournament of size 50, 100, 200, or 300 (chosen at random each time). The observation is always chosen from the last half of the permutation. The position-to-shift-in-front-of is chosen by a

tournament of the same size (each time) and is guaranteed to be at a location at least half way to the front of the permutation (starting at the 'deserving' observation). An observation that deserves to move earlier in the permutation determined by the following characteristics (in order):

- (a) unscheduled rather than scheduled
- (b) higher priority
- (c) later in the permutation

The position-to-shift-in-front-of has the opposite characteristics.

In the case of the genetic algorithms half of all children are created by a mutation and the other half by crossover. The crossover operator is Syswerda and Palmucci's position-based crossover (Syswerda & Palmucci 1991). Roughly half of the permutation positions are chosen at random (50% probability per position). These observation numbers in these positions are copied from the father to the child. The remaining observation numbers fill in the other permutation positions in the order they appear in the mother.

The final algorithm tested was HBSS with contention heuristics. HBSS does not use the permutation or the simple greedy scheduler. The observations are still scheduled one at a time, but next observation to schedule is chosen by a weighted roulette wheel. For a given observation, the windows are chosen by another weighted roulette wheel. Observation and window weights are assigned a dynamic weight which depends on the remaining observations and windows that do not violate any constraint except the duty cycle constraint. The duty cycle constraint is not considered by the contention heuristics for performance reasons. The weight of an observation is a function of how difficult it is to allocate a sensor and space on the SSR. The difficulty is the minimum of the contention any of an observation's windows faces in meeting the current constraints. The contention of each window is a function of the other observation's windows that have incompatible sensor use requirements or compete for the SSR. This is modulated by the need of each observation, which is a function of the observation's priority and number of windows currently believed not to violate constraints (remember that the duty cycle is not considered). Formally:

$$D_o = w_p P + w_s \min_w(C_s) + w_{ssr} \min_w(C_{ssr}) \quad (2)$$

$$C_s = \sum_{O_s} N_s \quad (3)$$

$$C_{ssr} = \sum_{O_{ssr}} N_{ssr} \quad (4)$$

$$N_s = P/|W_a| \quad (5)$$

$$N_{ssr} = SSR_a P/|W_a| \quad (6)$$

where D_o is an observation's difficulty, P is an observation's priority, \min_w is the minimum over all of an observation's windows that are believed not to violate

current constraints, C_s is a window's sensor contention, C_{ssr} is a window's SSR contention, $w_p = 1$, $w_s = 1$, $w_{ssr} = 1$, O_s is the set of observations that cannot use a window's sensor if the window is scheduled due to sensor use or slewing constraints, N_s is an observation's sensor need, O_{ssr} is the set of observations that use the SSR between the same two SSR dumps and are not believed to violate current constraints, N_{ssr} is an observation's SSR need, SSR_a is the amount of memory currently available in the SSR, and W_a is the set of an observation's windows are not believed to violate current constraints. All of these values, except the priority, are updated every time an observation is scheduled or found to violate the constraints; including the duty cycle constraint.

HBSS with contention heuristics is hundreds of times slower than the other techniques on this problem, and thousands of times slower on problems with much slower slewing motors (more windows are incompatible with each other). The contention heuristics require large, complex data structures to incrementally update the need, contention, difficulty and weight of each as-yet-unscheduled observation and its windows. This requires substantial memory. For example, this model problem with no limits on the slew angle could not fit in memory (there were more windows and more were incompatible with each other).

Experiment

To find the best algorithm for the model problem we compared a total of 14 techniques. These were ISAMP, HBSS, and every combination of three search techniques – hill climbing, simulated annealing, steady state GA, and generational GA – with three mutation operators – 1-15 random swaps, 1-15 temperature dependent swaps, and 1-15 squeaky shifts. Except for HBSS, 32 jobs with identical parameters (except the random number seed) were run for each algorithm. Each job generated approximately 100,000 schedules (the GA runs generated slightly more). These jobs took 2-3 hours each on one Athalon processor of our Linux cluster.

For HBSS, we ran the equivalent of eight 100,000 schedule jobs. Generating 100,000 HBSS schedules takes about 830 hours on the same Athalon processor (a couple hundred times slower than the other techniques). Due to time constraints, only 800,000 total HBSS schedules were generated, rather than 3.2 million for each of the other techniques. Thus, HBSS results are discussed separately below and are not included in tables and figures to avoid the impression the data are directly comparable. However, the HBSS results were so poor (barely better than ISAMP) that the difference in data size makes little difference to the conclusions.

Table 1 compares the mean fitness. Nearly all of the differences were statistically significant by both t-test and ks-test, with confidence levels usually far above 99%. We see that simulated annealing with temperature dependent swaps (SaTd) performs best with al-

gorithms using simulated annealing, hill climbing, temperature dependent swaps, and random swaps clearly leading. The message seems to be modify one schedule, take random steps, and restrict steps more and more as evolution proceeds.

Simulated annealing and hill-climbing with the squeaky shift operator are next. These outperform the genetic algorithm regardless of mutation operator. Again, within the genetic algorithms the squeaky shift mutator performs the worst.

ISAMP, as one might expect for random search, performs the worst. However, the very best HBSS schedule (out of 800,000 generated) was only a little bit better than the ISAMP mean.

The small standard deviations for all techniques suggests that all runs for a given technique get about the same fitness. Thus, even if the fitness landscape is multi-modal all the minima must be about the same. Figure 1, which shows the breadth of each fitness distribution over 32 runs, confirms this view. For this reason, we suspect that this problem requires mostly exploitation, rather than exploration, which also explains the poor GA results. Evolutionary change is spread out over the GA populations rather than concentrated on a single individual as for simulated annealing and hill climbing.

The squeaky shift mutators performance relative to random swaps suggests that it is smart in the wrong way. In preliminary experiments we also tried swapping, rather than shifting, observations and forcing observations to be swapped into certain parts of the permutation. The shift operator performed the best, but still not as well as the random swap mutator. See (Globus *et al.* 2003) for some of this data. If random outperforms intelligent, then clearly the intelligence is being applied in the wrong way. We do not understand the dynamics of permutation scheduling in any fundamental way, and we don't even know if the dynamics are fundamentally similar for different problems. Until a better understanding is reached, the random swap operators – with a decrease in the number of swaps as evolution proceeds – appears best.

Figures 2-4 show that the individual objectives in the weighted sum of equation 1 display much the same trend as the fitness. Simulated annealing and hill climbing with random swaps beats squeaky shifts and the genetic algorithm in almost every objective of the fitness function (figures 3-5). This is a bit surprising since the weights make the pointing and slewing objectives much less important than the sum of the priorities. However, notice that the range of average off-nadir pointing is very large suggesting that this measure made little difference, although the difference in means was almost always significant by both the t-test and ks-test.

Figure 5 indicates the number of unscheduled observations, an objective not found in the fitness function. Notice that the squeaky mutators in simulated annealing and hill climbing are worst than the genetic algorithms with random swaps. This suggests that the

squeaky mutators with simulated annealing and hill climbing do a better job of scheduling the high priority observations to make up for scheduling fewer observations.

Summary

We have compared 14 different techniques for scheduling EOS fleets on a realistically sized model problem. Simple techniques such as simulated annealing and hill climbing outperformed the genetic algorithm and HBSS with contention heuristics. Simple random swap mutation outperformed more 'intelligent' mutation. Reducing the number of random swaps as evolution proceeds also improves performance. The most 'intelligent' algorithm with no learning, HBSS with contention heuristics, barely outperformed random scheduling (ISAMP) in spite of requiring far more computing resources. Apparently, taking advantage of previous scheduling attempts, as simulated annealing, hill climbing, and the genetic algorithm does, has more value than large amounts of computation to choose just the right move to make. For this application and these techniques, fast and stupid with a little learning outperforms smart.

Acknowledgements

This work was funded by NASA's Computing, Information, & Communications Technology Program, Advanced Information Systems Technology Program (contract AIST-0042), and by the Intelligent Systems Program. Thanks to Alex Herz, Orbit Logic, Inc. and Jerald Arp, PhD., Manager of Technical Support, Space Imaging LLC for information regarding current EOS systems. Thanks also to Bonnie Klein for reviewing this paper and to Jennifer Dungan, Jeremy Frank, Robert Morris and David Smith for many helpful discussions. Finally, thanks to the developers of the excellent Colt open source libraries for high performance scientific and technical computing in Java (<http://hosc hek.home.cern.ch/hosc hek/colt>).

References

- Bresina, J. 1996. Heuristic-biased stochastic search. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence*.
- Frank, J.; Jonsson, A.; Morris, R.; and Smith, D. 2002. Planning and scheduling for fleets of earth observing satellites. In *Proceedings of the 6th International Symposium on Artificial Intelligence, Robotics, Automation and Space 2002*.
- Globus, A.; Crawford, J.; Lohn, J.; and Morris, R. 2002. Scheduling earth observing fleets using evolutionary algorithms: Problem description and approach. In *Proceedings of the 3rd International NASA Workshop on Planning and Scheduling for Space*.
- Globus, A.; Crawford, J.; Lohn, J.; and Pryor, A. 2003. Scheduling earth observing satellites with evo-

lutionary algorithms. In *Conference on Space Mission Challenges for Information Technology (SMC-IT)*.

Joslin, D. E., and Clements, D. P. 1999. Squeaky wheel optimization. *Journal of Artificial Intelligence Research* 10:353–373.

Potin, P. 1998. End-to-end planning approach for earth observation mission exploitation. In *SpaceOps 1998*.

Rao, J. D.; Soma, P.; and Padmashree, G. S. 1998. Multi-satellite scheduling system for leo satellite operations. In *SpaceOps 1998*.

Sherwood, R.; Govindjee, A.; Yan, D.; Rabideau, G.; Chien, S.; and Fukunaga, A. 1998. Using aspen to automate eo-1 activity planning. In *Proceedings of the 1998 IEEE Aerospace Conference*.

Syswerda, G., and Palmucci, J. 1991. The application of genetic algorithms to resource scheduling. In *Proceedings of the Fourth International Conference on Genetic Algorithms*, 502–508.

Search technique	mutation operator	abbreviation	mean fitness	fitness std. dev.
simulated annealing	temperature dependent swaps	SaTd	9205	20
hill climbing	temperature dependent swaps	HcTd	9310	21
simulated annealing	random swaps	SaSr	9311	19
hill climbing	random swaps	HcSr	9368	25
simulated annealing	squeaky swaps	SaSs	9489	19
hill climbing	squeaky swaps	HcSs	9507	24
generational GA	random swaps	GgSr	9700	38
steady state GA	random swaps	GsSr	9700	25
steady state GA	temperature dependent swaps	GsTd	9741	31
generational GA	temperature dependent swaps	GgTd	9834	24
generational GA	squeaky swaps	GgSs	9964	53
generational GA	squeaky swaps	GsSs	10010	46
ISAMP	random	ISAMP	10463	11

Table 1: Scheduling algorithms tested ordered by mean fitness. HBSS is left out since processing time did not permit 32 HBSS runs of 100,000 schedules each. The best HBSS fitness in the equivalent of 8 runs was 10438, a little better than the ISAMP mean but worse than all other techniques.




Figure 1: Comparing fitness for 32 runs. The boxes indicate the second and third quartiles. The line inside the box is the median and the whiskers are the extent of the data. Outliers represented by small circles.




Figure 2: Sum of the priority of unscheduled observations ($\sum_{I_u} P_o$ from equation 1).




Figure 3: Mean slewing time needed for each scheduled observation (M_s from equation 1).




Figure 4: Mean off-nadir pointing angle needed for each scheduled observation (M_p from equation 1).




Figure 5: Number of unscheduled observations ($|I_u|$ from equation 1).